

4. RELIABILITY

Seymour C. Himmel

Both NASA and the electric power industry are involved in the design, building, and operation of large and complex systems. These systems must operate efficiently and reliably. Failure can result in tremendous economic losses.

In the past, the electric power industry has been characterized by the application of modest incremental advances in technology as new plants were built. This relatively slow pace permitted careful and detailed investigation of the effects of the incorporation of these advances. In the process of preparing for this conference, it became evident that the electric power industry is on the threshold of the application of large advances in technology. This is quite similar to the situation in the late 1950's when, in order to undertake the space flight program, NASA had to take a large step in technology.

Such a step can, and in that case did, lead to difficulties. It is well known that in those days the space program was plagued by a lack of reliability in the systems being tried. To illustrate how bad the situation was, consider just the launch vehicles. In 1959, only 57 percent of the attempted launches succeeded. In 1960, only 59 percent met their flight objectives.

This was obviously a very unsatisfactory state of affairs. Upon close examination of the situation, it was found that the complexity of the systems, the new technology that was being used, and the rapid pace at which the systems were being developed all contributed to the poor performance obtained. Moreover, once the hardware was committed to flight, if something malfunctioned, it was not possible to get at it to find out what had happened - let alone to fix it. The electric power industry is fast approaching a similar situation. With some of the new reactors for power generation, one might as well be in space as it is frequently not possible to get at the hardware once a malfunction occurs.

In order to overcome the problems encountered, NASA developed a methodology for improving and maintaining the reliability and quality of the hardware. That the methodology is effective is illustrated by the fact that, in contrast to the less-than-60-percent success level previously noted, last year a 93-percent level was attained for the launching of twice as many vehicles. These vehicles, moreover, were of even larger size and greater complexity than those of 8 years ago.

The purpose of this paper is to acquaint you with this methodology. There are obviously many differences in our fields of endeavor, and all that is done to enhance the reliability of space systems may not be applicable or warranted in the power industry. But, it is believed that many of these techniques should be of value, and their application may permit the power industry to avoid some of the pitfalls encountered in the space program.

There are two basic parts to this methodology. The first is called reliability engineering and the second is called quality assurance. Where one ends and the other begins is a moot point. A somewhat arbitrary definition is used in this paper. What is important is that both are necessary.

In reliability engineering, it is established, by detailed examination and control of design, that a system is inherently reliable and, by test, that it can perform as desired. This process is called "qualifying" the system. In quality assurance, both inspection and testing are used to ensure that the system delivered by the manufacturer is as designed.

All of this sounds like an involved way of saying: "Use good engineering practice." In fact, that is all it is. It has been found, however, that today this is not the norm. Whether this results from the complexity of the systems, the rapid pace of development, the compartmentalization or specialization of engineers, or combinations thereof, is difficult to assess. The fact is that the problem does exist, and NASA has been forced to live with it.

RELIABILITY ENGINEERING

Reliability must be engineered into a system, it doesn't just happen. Reliability engineering consists of two basic parts, design and test. Some of the elements under these two major subdivisions are as follows:

(1) Design

- (a) Design criteria
- (b) Design practices
- (c) Design reviews
- (d) Reliability analyses
- (e) Failure mode analyses

(2) Testing

- (a) Breadboard
- (b) Prototype
- (c) Qualification of piece part, component, subsystem, system, and combined systems

Each of these elements plays an important role in the process of achieving a reliable system. All are required in order to attain a reliable system. Consider first the design of the system.

Design Criteria

After establishing what sort of system is required, that is, what it is to do and, conceptually, how it is to do it, the design criteria are established. Some of the criteria that must be established before undertaking a design are

- (1) Structural factors of safety
- (2) Electronic parts derating policy
- (3) Codes or standards
- (4) Functional specifications
- (5) Environmental specifications
- (6) Duty cycle
- (7) Electromagnetic interference

Many of these criteria are self-explanatory. Others are a little more involved, for example, electronic parts derating. Such parts normally have a service voltage and power specified by the manufacturer. The basis for the rating is somewhat obscure, and the physics of failure is not well understood. It is known that in solid-state devices, for example, there is a premium on small size of the part and that the failure mechanism is somehow related to operating temperature. By test, it has been established that, if the power level at which a part is used is cut by half, the frequency of random failure is reduced by a factor of 5 or better. Therefore, as a design criterion, it is stipulated that all such parts be used at no more than about 0.7 the specified voltage and current level. (The actual numbers vary with the type of part under consideration.) When it is recognized that a typical launch vehicle contains about 300 000 parts, most of them electronic, the product of all those one-fifth's in failure rate has great significance.

Environment plays a major role in how a system will perform. Vibration, shock, and temperature are easily envisaged. Space systems must also contend with the salt-water atmosphere at the launch complex in Florida and the potential for growth of fungus in a hot and humid environment.

Considerations such as these lead to what may be termed a design margin of safety. It establishes the degree of conservatism built into the system. Among the duties of reliability engineers is that of providing the data on which to base the design criteria.

Design Practices

Once the criteria are established, the design process can proceed. The reliability engineer is an important participant in system design. During all phases of the design process, he has the responsibility for ensuring that proper design practices are employed. This involves more than assuring adherence to the design criteria, it also includes, among others: (1) making sure that as few different varieties of parts as possible are used; (2) recognition of critical functions that would benefit from provision of redundancy in the system; and (3) making sure that, when a new technique or component is applied, there is a real need for it and that it is not just a manifestation of the designer's normal desire to do something new. In addition, he conducts a reliability analysis of the system. This analysis takes the statistical information on the failure rates of piece parts and projects the probability of failure of the systems components. Although the answers are expressed in probabilities, with lots of nines following the decimal point, the absolute numbers are not that important. What is important is the relative standing of the elements of the design. If an element is calculated to be significantly less reliable than most others, it is a candidate for redesign and/or redundancy. Redundancy, although frequently desirable and useful, is not a cure-all. Sometimes addition of redundancy can cause more trouble than it cures. Recognition of this when it occurs is the province of the reliability engineer. Examination of the design for proper design practice is a continuous part of the design process.

Design Reviews

Another important element of the reliability engineering process is the design review. Design reviews are normally conducted at three points during the design process. The first one is the conceptual design review, which occurs relatively early in the cycle. At this point, the basic elements of the system are selected, and the mechanization schemes to satisfy the fundamental requirements are established. Preliminary design follows during which parts are selected, specifications are drafted, and preliminary drawings are made.

At this point, the intermediate design review takes place. During this process all the design data and schematics are reviewed to establish that the system can do the required job and that the components or parts selected satisfy the design criteria. When this has been established, a very important step begins. This is the failure mode and effects analysis of the proposed design. The design is examined to answer the following questions:

- (1) How can a part or component fail?
- (2) For each mode of failure - what effect does it have on system performance?
- (3) How critical is this effect?
- (4) Can the problem be obviated?

Basically, the failure mode analysis asks "what if?" and attempts to answer the question. To illustrate the sort of thinking involved, consider the schematic of a launch vehicle hydraulic system shown in figure 4-1. As can be seen, it is basically a simple system. A pump takes hydraulic fluid from a reservoir, raises its pressure, and passes it through a check valve to the engine gimbal actuator, whence it returns to the reservoir. The system pressures are measured by three transducers, two on the high-pressure side and the other on the return side. During the failure mode analysis, the question was asked: "What if a high-pressure transducer fails?" The transducer that had been selected was examined. It is shown in figure 4-2. It consists of a spirally wound Bourdon tube that moves a potentiometer. The working pressure of the system is 3000 psi; the gage was rated to 0 to 3500 psi. It seemed as if this was sufficient margin. The weakest point in the Bourdon tube is where it is brazed to the inlet fitting. If this were to break, the case of the transducer would be subjected to full line pressure and, for the instrument that had been selected, would fail because the electrical connector was merely soldered on to the case. In a short time all the supply of fluid would be exhausted overboard, and control of the vehicle would be lost. This possibility was, of course, totally unacceptable and was avoided by selecting a transducer case that could withstand full line pressure in the event of a failure of the sensing element. It was tolerable to lose a measurement, but not the whole vehicle.

This is not a hypothetical situation. It actually existed in a vehicle design and probably would have gone unnoticed if the design had not been subjected to a failure mode analysis.

As a result of such work at all levels of the design, the system design is modified to provide inherent reliability. One other important element is provided from the results of such analyses. That is, as each potential failure mode is identified, the test requirements for acceptance testing are established. These should be selected so that each of the modes can be detected. Thus, some of the instrumentation requirements for the system are established.

Following the detailed design of the system, everything is examined again in the final design review. This repetition is required because, at this juncture, the physical arrangement of the hardware has been established and this may have altered the intent of the design as established during the intermediate design review.

Figure 4-3 illustrates the sort of problem that can arise. Again, this actually occurred. It was decided to provide redundant gates in a particular circuit of an

autopilot because of its criticality. The schematic of the design is shown in the left of the figure. As can be seen, if one of the two gates fails, the signal can still get through - just what was wanted. When this was translated into a physical detailed design, however, the situation shown by the solid lines in the right sketch resulted. It looks, schematically, just like what was wanted - except that these gates are part of a single integrated circuit chip and are thus not truly independent. Were the chip to fail, both gates would most probably fail. The desired redundancy was therefore not attained. The solution is shown by the dashed lines. In this case, the second gate was provided on another integrated circuit chip, and true redundancy was obtained. It is to uncover deficiencies like this that final design reviews are held.

As may be inferred from the examples given, it takes a special "breed of cat" to be a good reliability engineer. He must be a highly competent engineer, versed in design analysis, knowledgeable of how things are done in the "real world," and have the patience of Job. When such a paragon has been found, there is a most important thing that should not be done with him. That is, he must not be put to work for the project manager. He must not be subject to the pressures on the project to meet costs and schedules. He must be free to "nit pick" and to act as the conscience of the project. In NASA, the reliability engineer works with the project manager but out of an independent office that reports to the Center director.

Testing

After the final design review and the resultant design modifications, the first item can be built. Despite all the effort described, not a thing has been built; that is, nothing resembling the final design has been made. Frequently, as part of the design process, a breadboard setup is made. This is really a design tool - an analog simulation of the schematic to determine, for example, whether a complex circuit will function as intended.

The first item that resembles the flight item is the engineering model or prototype, this is normally built in an experimental shop. With it the first step in proving the reliability of the system is taken. The testing that is conducted is called design evaluation testing and is very similar to the formal qualification program. The nature of the testing is as follows:

- (1) Functional tests
 - (a) Design condition
 - (b) Specification range, off standard
 - (c) Tests for marginality

(2) Environmental tests

- (a) Pressure
 - (b) Thermal
 - (c) Vibration
 - (d) Shock
 - (e) Atmosphere
 - (f) Electromagnetic interference
 - (g) Rain
 - (h) Sand
 - (i) Fungus, etc.
- } combined where possible

(3) Stress limit tests

(4) Life tests

- (a) Specified life
- (b) Extended life

It is important to note that the applicable portions of the tests listed are conducted at various levels of the hardware, that is, first on the piece parts, then on the components or modules, next on subsystems, and, finally, at the system level. Ultimately, if at all possible, similar tests are conducted at the combined systems level, that is, all-up with all systems interacting. The tests are described briefly in the following paragraphs.

First come the functional tests. These tests prove that the system can work as designed and are conducted at both nominal conditions and over the range of specified operating conditions, that is, over the range of inputs for which the system should perform properly. After this, the variables are permitted to exceed the specifications to determine whether a marginal condition exists; that is, whether an input slightly out of specification will cause the system to fail or malfunction.

When this series is completed, the environmental testing phase begins. The system is subjected, while operating, to the types of environments that it will experience in service. In qualification testing, the specimen is normally subjected to environmental conditions more severe than it would be expected to encounter in service. For example, in vibration testing it is normally required that the design be capable of surviving at least $1\frac{1}{2}$ times the expected flight vibration levels.

If it is at all possible, the environmental tests are conducted under combined environments such as the system will actually experience. On a launch vehicle, for example, a subsystem, or "black box," will be subjected simultaneously to vibration, thermal changes, and ambient pressure changes. The tests are designed so that these factors are applied as they would occur in flight.

If something fails during such tests, design changes are made to preclude such failure. Appropriate portions of the test program are repeated to prove the "fix"

before proceeding to other parts of the program.

When the environmental tests have been completed successfully, the subsystem enters the stress limit test. In this test, the input levels are incrementally increased. For example, the temperature level at which a "black box" is operated is extended on both high and low sides, to ascertain at what temperature it malfunctions. This type of test determines any critical weakness or marginality in the design and establishes the actually achieved design margins.

The last type of test listed, the life test, consists of repetitive cycles of operation at the anticipated environment. This is really a sort of fatigue test during which the time it takes for a system to "wear out" is ascertained.

When the prototype and all of its parts have successfully completed the tests with all the modifications incorporated, the next step can begin; that is, the drawings and process specifications are released to the production department for manufacture. The first few items produced in these shops are then subjected to the whole gamut of tests, this time for formal qualification. Recall that the prototype hardware was built in an experimental shop and that there is a vast difference between it and a production shop. This series of tests on the production hardware is viewed as qualifying both the hardware built in the production shop and the shop itself. When all these tests have been successfully completed, the design is considered to be qualified for flight.

QUALITY ASSURANCE

After all of the effort that went into the design and qualification process just described, it would seem that the problems are over and all that has to be done is to build or buy the product of all this work. Unfortunately, this is not the case. The system must be policed to ensure that the product will be precisely as designed. This is the function of the quality assurance program. The tools of this discipline include specifications, process control, inspection, acceptance testing at all levels, a closed-loop failure analysis and corrective-action system, and stock controls. In order to get a high-quality product each of these tools must be employed in a rigorous manner.

Specifications

First, consider the specifications that govern the items that are built or purchased. It is important that these be written as completely and precisely as pos-

sible. They should contain negative as well as positive statements. It must be remembered that you get only what you contract for.

The following example shows how an incomplete specification caused much trouble. In an autopilot, a large number of a particular type of transistor were employed. It had been a long, hard task to qualify a particular vendor's product for this application, and the systems were working wonderfully. Suddenly, the autopilot began to experience unexplainable functional failures during testing. Nothing had been changed; all parts were being purchased to the specifications always used, and the manufacturing processes employed in the electronics shop were the same and were being followed carefully. After extensive testing, all elements had been eliminated as the source of the problem, except this type of transistor. After some investigation it was found that, without notification, the vendor had "improved" his product. Figure 4-4 shows the gain characteristics of the transistor. The solid curve shows the gain as a function of the collector current of the original part. The dashed curve shows the characteristics after the improvement. The difference is apparent.

The purchase specification required that the gain at a specific current lie in the range enclosed by the circle. It did not say that the gain at lower currents should follow the solid curve. This had been considered an implicit characteristic of the device. Because many of the transistors in the autopilot actually operated at a much lower current (as shown by the dashed vertical line), the change of gain by a factor of 2 drew the circuits into instability. The new model transistor still met the specifications as written, and the vendor, who did not know the application, had no idea of the trouble a normal product improvement would cause. The specifications now state "No changes without telling us."

Inspection

Inspection is a vital element of quality assurance. In the qualification tests, not only is a design being qualified, but the manufacturer and the manufacturing process are also being qualified. The item must always be made in the same fashion in which the one that was qualified was made. This requires that inspection be thorough, that the most critical items are identified for special attention, and that strict adherence to all process specifications be maintained.

At one time, an inspector was a highly skilled, experienced, and capable craftsman who could perform all the tasks he was called on to inspect. Unfortunately, this is no longer generally true. Inspectors frequently do not understand what they are looking for. They merely follow instructions, and these, perforce, cannot completely describe what constitutes a good piece of work.

Testing

For this reason, reliance must be placed on the most powerful tool at our disposal - testing. Everything used is subjected to a rigorous acceptance test program at all hardware levels. The testing encompasses both functional and environmental tests. In fact, the tests are very similar to their counterparts during qualification testing. The principal difference is that the environmental levels are normally limited to what the expected service levels will be. The value of acceptance testing under appropriate environmental conditions cannot be overemphasized. Again, an example will illustrate the point.

During the acceptance flow response test, an hydraulic pump was rejected as being excessively noisy and exhibiting abnormal vibration. The pump is a conventional constant-pressure, positive-displacement, variable-flow device. A sample of hydraulic oil from the pump-case drain port contained a large number of bronze-colored particles and several pieces of broken ball-bearing retainer. Upon disassembly, it was found that the pump cylinder block bearing had been installed backward with the thrust face away from the bearing retainer. A sketch of the bearing is shown in the left of figure 4-5. Note the asymmetry of the bearing. The pump vendor's assembly instructions stated "Assemble with face marked THRUST next to bearing retainer." This marking is shown in the right of the figure. The outer race can support only minimal thrust in the reverse direction. When the pump was under test, the bearing broke. So much for the quality of the inspection. As this pump is not run until the rocket engine turbopump drives it at engine ignition and the pump does not fail immediately, it would have done so during flight and would have caused a mission failure. Only the acceptance tests prevented the improperly assembled pump from being installed on a vehicle.

Failure Analysis and Corrective Action

A necessary concomitant of any testing is a closed-loop failure analysis and corrective-action system. Unfortunately, design and manufacturing techniques are not infallible, and despite all the effort applied to preclude it, a subtle design deficiency may sneak through and cause trouble. When something malfunctions in a test, it is not enough to fix the one item and then proceed. It is vital that the cause of the failure be identified and corrective action be taken on all such devices. A subtle design deficiency and the failure analysis and corrective action taken are illustrated in figure 4-6.

This figure depicts an electrohydraulic servovalve that controls the engine

gimbal actuators. Basically, it is a differential electric motor driving a pilot valve that ports hydraulic oil to either side of a working cylinder to move the engine. The three leads shown carry the current for moving the electric motor in response to commands from the autopilot servoamplifier.

This valve had been subjected to a thorough qualification program and had come through with flying colors. During acceptance testing of a batch of these valves, one of them failed during the vibration test. The failure was a short circuit in one of the three wires to the motor. What had happened is shown in the photograph in figure 4-6. Notice how the three wires are pinched between the pole piece of the motor and the case. Upon investigation, it was found that the gap between the case and pole piece through which the wires from the armature had come was not quite big enough for a single wire. Thus, when the motor was assembled, the wires were slightly pinched. The amount of pinching involved was not enough to break through the insulation, however. The insulation is Teflon, and this material is subject to cold flow. When the valve was assembled at the vendor's plant, it successfully passed his tests. The pinched insulation, however, continued to flow. During the acceptance test, the added vibration (which is like that which would be encountered in flight) completed the insulation flow, and the short occurred. In flight, such a failure would have been catastrophic.

The failure analysis showed that, in the process of assembling the valve, the motor leads are soldered to a bulkhead connector and then the connector and wires are pushed into the case and the connector is screwed in place. Unless care is taken during this step, it is easy for the wires to slip out of the proper feedthrough area and be pinched over the sharp edge of the pole pieces. A number of steps were taken to correct this situation. First, the edge of the pole piece, which had been square and relatively sharp, was given a generous radius. And second, special assembly instructions and inspection callouts were implemented.

The most important aspect of this incident is the fact that a failure analysis was conducted. It would have been very easy to simply reject the valve and send it back to the manufacturer. Such action would, however, have left the valve with a large potential for catastrophic failure in a future flight.

Most of the preceding discussion and examples have been concerned with problems within a single component or "black box." Another important factor to consider is that of interactions or interfaces among systems. Whenever there is an interface, there is a potential for an incompatibility that can lead to trouble. Figure 4-7 illustrates a problem of this nature. This example demonstrates the importance of failure analysis as well as the pitfalls of an interface.

The figure shows a schematic of a portion of the propellant utilization system of a launch vehicle. The function of this system is to regulate the ratio of fuel and

oxidizer flows to the rocket engine so that the supply of both propellants is depleted at the same time. The portion of the system circuit shown is that between the error detector and the fuel valve controller. The solid lines are the original configuration. This is basically a very simple circuit. The information tap provided a means for measuring the output of the error detector. The signal went to both the telemetry system on the vehicle and through a landline to a recording instrument in the blockhouse. It was noticed during prelaunch testing that, when the umbilicals were ejected, a shift in the error signal was indicated by the telemetry signal. Although the shift was relatively minor, it should not have occurred. A failure analysis was instituted. A detailed analysis of the circuit at the launch site showed that, when the umbilicals were connected, the instrumentation in the blockhouse was loading the circuit and introducing a bias which evidenced itself when the umbilical connection was broken. As the vehicle was checked out in supposedly the same way in the factory, the obvious question was "Why didn't this show up then?" The answer was, quite simply, that the combined impedance of the blockhouse instrument and long line at the launch site was much different from that of the equipment used in the factory. The factory gear did not load the circuit. This is an almost classic example of the importance of using the same type of test gear in both the factory and field - an interface problem. Thus, the problem of the error signal shift was resolved.

While the problem was being investigated, it became obvious that, if there were a failure in the instrumentation system, it could cause a failure of the propellant utilization system. The original instrumentation line went directly to the telemetry set. If a short developed in the telemeter, the output of the error signal would be shorted, and control of the fuel valve would be lost. Thus, the desire or need for information of system performance parameters resulted in a design that introduced an unacceptable failure mode. Fortunately, the solution for both problems was straightforward and is shown by the small-dashed lines. The error signal was picked up as shown and fed through isolating resistors. The resistors serve two purposes. First, they provide a high impedance ratio so that the instruments can no longer load the circuit. Second, they isolate the instrumentation system from the propellant utilization system so that, if there is a short in the information system, it will not destroy the functional system.

This concludes our brief tour through a chamber of horrors. These are but a few examples of the sort of problems that the methodology NASA has evolved has helped prevent or catch. Although most of the examples have been of an electrical nature, a similar set concerned with purely mechanical components could have been used.

Operational Testing

The reliability and quality assurance program does not end with getting flight hardware through the acceptance test at the factory at the subsystem level. It continues all the way to launch and throughout the flight.

After the subsystems have been through their individual acceptance tests, they are assembled onto the vehicle in the factory and each system is tested. When the systems tests are completed satisfactorily, a combined systems test is conducted with all systems being required to operate while interacting with each other as they would in flight. (Ideally, this should be done in a simulated flight environment.) Only when this test has been passed is the vehicle ready to ship to the launch site.

At the launch site, the systems level tests are repeated to determine whether any damage or degradation has occurred during transportation and to prove that the airborne hardware is compatible with the ground-support equipment. Then a series of composite tests is undertaken to again demonstrate the integrity of the vehicle. Typically, there are three composite tests during which everything that will happen during launch and flight, short of actually igniting the engines, is checked out.

When all this testing is over, a data and readiness review is held. Each box, or major component, has a "history jacket" which documents its life story during manufacturing and testing. The data from all tests are now reviewed to ascertain whether each system is, and has been, performing within specifications and whether the data exhibit any trends indicative of deterioration - even if all readings are within specifications. Only when each system proves to be acceptable on the basis of all testing, is it considered ready for flight.

Subsequent to flight, all data from the flight are evaluated to ascertain whether the systems performed as predicted and whether any anomalies occurred.

CONCLUDING REMARKS

The reliability and quality assurance program described briefly herein is, however, no panacea. It can lead to a paper blizzard because all the things that are discovered must be documented and communicated in the hope of precluding a recurrence. This can lead to a reliance on paper rather than on sound engineering evaluation and judgement. The program does help to reduce the frequency of undetected human errors but, alas, does not eliminate them. Also, it takes money, sometimes a lot of it, but it can save much more - and that makes it worthwhile.

What has been learned in the struggle to achieve reliability?

First, testing is the most potent and least expensive aspect of the program

when measured against the cost of a flight failure. The equipment and facilities required to perform full-systems environmental tests are worth the cost. For a while, the ability to test the hardware of both spacecraft and launch vehicle systems was comparable. Now, spacecraft can be more effectively tested because their size is compatible with existing environmental facilities. It is now standard practice to put the whole spacecraft into an environmental chamber wherein almost the total flight environment can be simulated. We look forward to the day when the launch vehicles may be tested in the same manner.

Second, designs must possess adequate margins of safety. Getting such margins must be a conscious part of the design effort.

Third, testability must be designed into the hardware. Also, the test procedures and test equipment are part of the design process - they are just as important in achieving reliability as the functional design of the flight hardware.

Finally, it is essential to police the manufacturing and test processes. Close adherence to procedures is mandatory. Parts should be standardized as much as possible, and good housekeeping practices must be followed at all times and places.

As stated at the beginning, all that NASA does to achieve reliability may not be applicable to the electric power industry. However, the two fields have a basic similarity. Both are engaged in activities that are characterized by small build rates of the ultimate end item - be it a powerplant or a launch vehicle. But there are many components and subsystems that are common to many installations. Perhaps group action within the industry in the design and specifications of such common elements would be an effective way of reducing the cost to an individual company of implementing reliability and quality assurance requirements on purchased items.

High reliability has been difficult to attain, just as difficult to maintain, and even more difficult to improve. It is hoped that the experience of NASA will in some way be of value to the electric power industry.

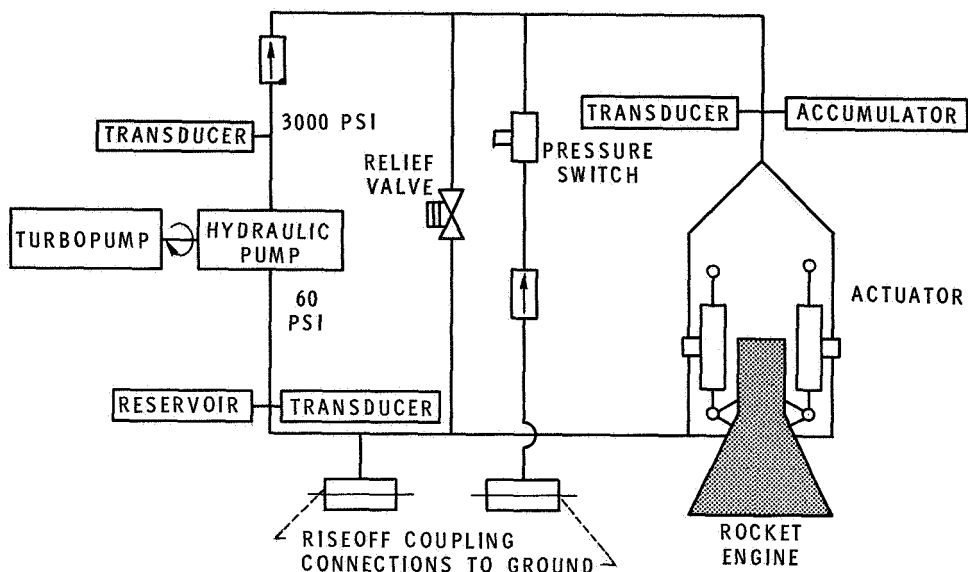


Figure 4-1. - Launch vehicle hydraulic system.




Figure 4-2. - Pressure transducer, 0 to 3500 psia.

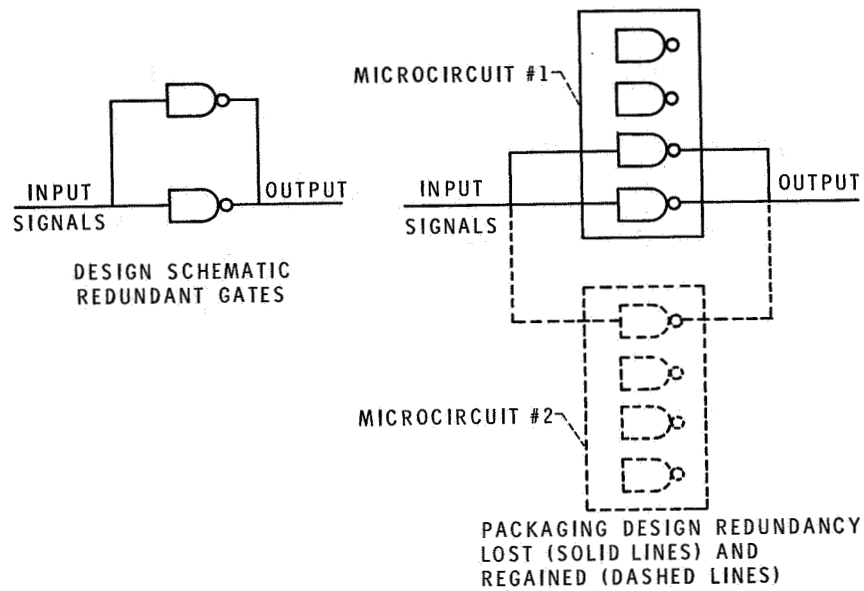


Figure 4-3. - Schematic of autopilot circuit gating.




Figure 4-4. - Transistor characteristics.




Figure 4-5. - Hydraulic pump bearing. Arrows show correct direction for loading races.




Figure 4-6. - Electrohydraulic servovalve for engine gimbal control.




Figure 4-7. - Schematic of portion of launch vehicle propellant utilization system.